Per Zone Variable BPI for Improving Storage Device Capacity and Yield

By

Saeed Asgari

Mathew Vea

And

George Iszlai

Field of the Invention

The present invention relates generally to the storage of information on fixed storage media, and more particularly to improving storage of information on rotating magnetic media such as disks in a disk drive.

Background of the Invention

Data storage devices such as disk drives are used in many data processing systems for data storage. Typically a disk drive includes a magnetic data disk having recording surfaces with concentric data tracks, and a transducer head paired with each recording surface, for writing data to, and reading data from, the data tracks. Each paired magnetic head and media surface couples to provide a unique data recording capability which depends on the fly height of the head from the recording surface, the quality/distribution of magnetic media on the recording surface, and the magnetic properties of the magnetic head.

Conventional methods of recording data using the paired head and recording surface are inefficient because they do not take into consideration the differences in data recording capabilities between one pair of head and recording surface, and another head and surface pair. Though the heads are designed to perform identically in read/write operations, in practice different heads in a disk drive can have different read/write performance capabilities. Lower performing heads cannot read/write data as that of other heads in the disk drive. Typically, a single error rate level and a single storage capacity level are used to record data for all the pair heads and surfaces. This results in inefficient data storage for

those pairs of heads and surfaces that can store more data. It also lowers the qualification yields of the disk drives because one or more pairs of heads and surfaces do not record data at the qualifying error rate and capacity levels.

Further, in high data rate design of disk drives, as the recording density (i.e. bits-per-inch and/or tracks-per-inch) is increased, maintaining transducer head tolerances has become a challenge. Variance in the relative head performance distribution increases with increasing data density. In conventional disk drives, the drive yield and capacity suffers as a result of head performance variations in disk drives.

One method of increasing the data storage capacity of a disk drive includes increasing the areal density of the data stored on the media surfaces (bits/sq. in. -- BPSI). Areal density is the track density which is the number of tracks per radial inch (TPI) that can be packed onto the media/recording surface, multiplied by the linear density (BPI) which is the number of bits of data that can be stored per linear inch.

Conventional processes for qualifying disk drives scrap a disk drive when the measured disk capacity of the disk drive is less than a target disk capacity. Conventionally, each recording surface is formatted to store the same amount of data as every other recording surface. Thus, a recording surface that has a low error rate is formatted to the same TPI and BPI levels, as a recording surface having a high error rate, even though it can store more data. However, by adopting a single TPI and BPI level for every recording surface, conventional processed fail to account for the differences in sensitivity and accuracy of the paired head and recording surface, which results in less data storage and more waste of space on each recording surface. This also results in lower overall yields of disk drives because if even a few of the recording surfaces do not meet their targeted capacity, the sum of the surface capacities of all the media surfaces will be less than the target capacity, causing the entire disk drive to fail.

U.S. Patent Nos. 6,091,559 and 5,596,458 provide for recording at different BPI on different recording surfaces, however, such methods do not take into consideration multiple constraints, including head performance across the stroke per disk surface affecting disk drive capacity, disk drive performance requirements (e.g., throughput) and manufacturing requirements (e.g., test time). Zone frequencies are selected based on measurement of a single metric on one head.

There is, therefore, a need for a method of storing data in a disk drive which improves disk drive yield while meeting the desired target drive capacity or increasing the drive capacity while meeting a desired drive yield by taking advantage of the head performance variation.

Summary of the Invention

The present invention satisfies these needs. According to one embodiment of the present invention, a population of disk drives is selected, and head performance measurements are taken for each selected media surface location at different frequencies. Performance distributions are obtained from the measured data, and a format optimizer uses the distributions to obtain a design of different frequencies across the media surface zones, and determine head allocation. Once the different frequencies for the zones have been determined, then in each disk drive, the heads are assigned to the predetermined frequencies optimized for. As such, the present invention allows maintaining consistent performance (both sequential and random throughput) across a population of disk drives, and reduced test time. This is accomplished by determining head performance and design of format at development/design time, and assignment of heads to different frequencies at manufacturing time. Therefore, predetermined design of formats is performed off-line, and then marries to a manufacturing test process for assignment of heads to different frequencies.

In one example, the density/format for each recording surface zone and the number of heads allocated to each density, are preselected at design time, and at manufacturing time heads are assigned to higher/lower density formats. Unlike conventional methods, head allocation and assignment is per head per zone, taking into consideration head performance variation across zones. As such, if a first head that performs well at ID but poorly at OD, and a second head has reverse performance, that performance is traded off wherein the first head is assigned to high density at ID and at low density at OD, and the second head in the opposite fashion. In the per zone variable BPI for improving capacity according to the present invention, several manufacturing and customer constraints are taken into consideration. Performance of each head across the stroke, as well performance variation from one head to another, is utilized in designing the density format and assignments of heads to the density formats.

The present invention provides a variable BPI storage format as a function of storage zones in storage devices, such as disk drives, based on transducer head performance variations between different heads in a set of disk drives. The present invention provides a method of defining such a storage format in multiple data storage devices, each data storage device having a plurality of storage media and a plurality of corresponding data transducer heads, each transducer head for recording on and playback of information from a corresponding storage medium in multiple zones, wherein each zone includes a plurality of concentric tracks for recording on and playback of information. The method includes the steps of: selecting a plurality of said data storage devices; for each selected data storage device, measuring a record/playback performance capability of each head at one or more read/write frequencies per zone; based on said performance capability measurements, generating storage density distributions corresponding to at least a number of the heads in said selected data storage devices; selecting a group of read/write frequencies for said multiple data storage devices, two or more frequencies for each zone, based on said storage density distributions; and

thereafter, during manufacturing, assigning one of said read/write frequencies to each head based on performance capability of that head per storage device.

Brief Description of the Drawings

These and other features, aspects and advantages of the present invention will become understood with reference to the following description, appended claims and accompanying figures where:

FIG. 1A shows an example partial schematic diagram of a disk drive with an example data storage format according to the present invention;

FIG. 1B shows another example schematic of the disk drive of FIG. 1A illustrating disk drive electronics;

FIG. 1C shows an example surface format for data storage according to the present invention;

FIG. 1D shows an example diagram representing the general zone layout of a disk drive with N disks and 2N heads, depicting different heads in a section of a zone on different disk surfaces;

FIG. 1E shows another example of capacity zone layout on a surface a disk;

FIG. 1F shows example of a series of radial zones on a disk surface, wherein each zone includes multiple virtual cylinders;

FIG. 1G shows an example representative data track layout for each of several virtual cylinders in a zone on different disk surfaces with corresponding heads;

FIG. 1H shows another example servo track and data track layout for a zone on different disk surfaces with corresponding heads, wherein the number of servo tracks and data tracks in different virtual cylinders of a zone on different disk surfaces are the same;

FIG. 1I shows another example layout wherein the servo and data track layout varies from zone to zone on a disk surface;

FIG. 2A shows an example flow/functional diagram of embodiment of steps of generating the layout of FIG. 1I according to the present invention;

FIG. 2B shows a graph of playback error measurement for a head at a zone at different recording frequencies;

FIG. 2C shows an example joint BPI distribution plot;

FIG. 2D shows an example histogram of the frequency capabilities of the heads in a set of disk drives at a zone at a fixed target error rate;

FIG. 3 shows an example flowchart of an embodiment of steps of variable zoning data collection process of FIG. 2A;

FIG. 4 shows an example flowchart of an embodiment of steps of vertical zoning post-measurement processing and per zone BPI distribution extraction process of FIG. 2A;

FIG. 5 shows an example flowchart of an embodiment of steps of vertical zone assignment process of FIG. 2A; and

FIG. 6 shows an example flowchart of an embodiment of steps of format generation and optimization interaction process of FIG. 2A.

Detailed Description of the Invention

Data storage devices used to store data for computer systems include, for example, hard disk drives, floppy disk drives, tape drives, optical and magneto-optical drives, and compact disk drives. Although the present invention is illustrated by way of an exemplary magnetic hard disk drive 100, the invention can be used in other storage media and drives, including non-magnetic storage media, as apparent to one of ordinary skill in the art and without deviating from the scope of the present invention.

Referring to FIGs. 1A-C, an exemplary hard disk drive 100 is diagrammatically depicted for storing user data and/or operating instructions for a computer system 54. The hard disk drive 100 comprises an electro-mechanical head-disk assembly 10 shown in FIG. 1A as including one or more rotating data storage disks 12 mounted in a stacked, spaced-apart relationship upon a rotating spindle 13. The spindle 13 is rotated by a spindle motor 14 at a predetermined angular velocity.

Each disk 12 defines at least one media surface 23, an usually two media surface 23 on opposing side of each disk 12. Each media surface 23 is coated is coated with magnetic or other media for recording data. The spindle drive motor 14 turns the spindle 13 in order to move the disks 12 past magnetic transducer heads 16 suspended by suspension arms 17 over each media surface 23. Generally, each magnetic head 16 is attached to the suspension arm 17 by a head gimbal assembly (not shown) that enables magnetic head 16 to swivel to conform to the media surfaces on the disks 12. The suspension arms 17 extend radially from a rotary voice coil actuator (not shown). An actuator motor 54 rotates the actuator and head arms and thereby positions the magnetic heads 16 over the appropriate areas of the media surfaces 23 in order to locate and read or write data from or to the storage surfaces 23. Because the disks 12 rotate at relatively high speed, the magnetic heads 16 ride over the media surface 23 on a cushion of air (air bearing). Each magnetic head 16 comprises a read element (not shown) for reading magnetic data on magnetic storage media surfaces 23 and a write element (not shown) for writing data on the media surfaces 23. Most preferably, although not necessarily, the write element is inductive and has an electrical writing width which is wider than an electrical reading width of the read element, which is preferably of magnetoresistive or giant magnetoresistive material.

Referring to FIG. 1C, each media surface 23 is divided into a plurality of concentric circular tracks 30 that each have individually addressable portions 35, such as sectors, in which data is stored in the form of magnetic bits. The data sectors 35 are separated by embedded narrow servo sectors or spokes 25 which include a series of phase-coherent digital fields followed by a series of constant frequency servo bursts. The servo bursts are radially offset and circumferrentially sequential, and are provided in sufficient numbers such that fractional amplitude signals picked up by the read element from portions of at least two bursts passing under the read element enable the controller 57 to

determine and maintain proper head position relative to a data track 30. One example of a servo burst pattern for use with an inductive write element/magneto-resistive read element head 16 is provided by commonly assigned U.S. Patent No 5,587,850, entitled: "Data Track Pattern Including Embedded Servo Sectors for Magneto-Resistive Read/Inductive Write Head Structure for a Disk Drive", incorporated herein by reference.

The drive controller 57 controls operation of the pairs of magnetic heads 16 and media surfaces 23 to read and write data onto each media surface 23. The drive controller 57 preferably comprises an application specific integrated circuits chip which is connected by a printed circuit board 50 with other chips, such as a read/write channel chip 51, a motors drive chip 53, and a cache buffer chip 55, into an electronic circuit as shown in FIG. 1B. The controller 57 preferably includes an interface 59 which connects to the host computer 54 via a known bus structure 52, such as ATA or SCSI.

The controller 57 executes embedded or system software comprising programming code that monitors and operates the controller system and driver 100. During a read or data retrieval operation, the computer system 54 determined the "address" where the data is located on the disk drive 100, i.e., magnetic head number, the track 30 , and the relevant portion(s) 35 of the track 30. This data is transferred to the drive controller 57 which maps the address to the physical location in the drive, and in response to reading the servo information in the servo sectors, operates the actuator motor 54 and suspension arm 17 to position a magnetic head 16 over the corresponding track 30. As the media surface 23 rotates, the magnetic head 16 reads the servo information embedded in each spoke 25 and also reads an address of each portion 35 in the track 30. When the identified portion 35 appears under the magnetic head 16, the entire contents of the portion 35 containing the desired data are read. In reading data from the media surface 23, the read element (not shown) senses a variation in an electrical current flowing through a magnetoresistive sensor of the

read element (not shown) when it passes over an area of flux reversal on the surface 23 of the media. The flux reversals are transformed into recovered data by the read/write channel chip 51 in accordance with a channel algorithm such as partial response, maximum likelihood (PRML). The recovered data is then read into the cache memory chip 55 of the disk drive 100 from whence it is transferred to the computer system 54. The read/write channel 51 most preferably includes a quality monitor function which enables measurement of the quality of recovered data and thereby provides an indication of data error rate. One channel implementation which employs channel error metrics is described in commonly assigned U.S. Patent No., 5,521,945 to Knudson, entitled: "Reduced Complexity EPR4 Post-Processor for Sampled Data Detection", incorporated herein by reference. The indication of recovered data error is used in order to select linear data density, track density and/or error correction code levels, in accordance with principles of the present invention, as more fully explained hereinbelow.

Writing or storing data on the media surface 23 is the reverse of the process for reading data. During a write operation, the host computer system 54 remembers the addresses for each file on the media surface 23 and which portions 35 are available for new data. The drive controller 57 operates the actuator motor 54 in response to the servo information read back from the embedded servo sector 25 in order to position a magnetic Head1, settles the head 16 into a writing position, and waits for the appropriate portions 35 to rotate under the head 16 to perform the actual writing of data. To write data on the media surface 23, an electrical current is passed through a write coil in the inductive write element (not shown) of the head 16 to create a magnetic field across a magnetic gap in a pair of write poles that magnetizes the magnetic storage media coating the media surface 23 under the head 16. When the track 30 is full, the drive controller 57 moves the magnetic head 16 to the next available track 30 with sufficient contiguous space for writing of data. If still more track capacity is required, another head 16 is used to write data to a portion 35 of another track 30 on another media surface 23.

In one aspect, the present invention increases the data storage capacity and yield of data storage devices having a plurality of media surfaces 23, such as hard disk drive 100 including disks 12 covered with magnetic media.

Overview of general method vertical zoning

In every disk drive, there is a distribution associated with head/media pair performance in that disk drive. The present invention takes advantage of that distribution to determine different linear density recording frequency assignment for heads, and optionally track allocation. According to one embodiment of the present invention, a set of disk drives is selected, and head performance measurements are taken for each selected media surface location in the disk drives at different frequencies.

Empirical frequency capability histograms are extracted at a given known target performance metric from measurement data. Probability cumulative distribution functions (such as joint probability distributions) are estimated from the histograms and fed into a format optimizer to obtain and design (vertically zoned) frequency format profiles (i.e., across the stroke and media surface zones) as well as optimal number of head allocations to frequencies. Once frequency format profiles and optimal number of head allocations are designed and pre-determined, during a test process, every head at every zone is assigned to one of the multiple pre-determined frequencies based on the head's performance capability.

As such, the present invention allows maintaining consistent performance (sequential/random throughput) across several of disk drives, without introducing significant additional test time. This is accomplished by determining head performance and design of format at development/design time, and assignment of heads to different frequencies at manufacturing time. Therefore, the predetermined design of frequency format profiles and (optimal) number of head

allocations are performed off-line while the assignment of heads to different frequencies is performed during the test process such as during manufacturing.

Unlike conventional methods, in the present invention head allocation and assignment is per head per zone, taking into consideration head performance variation across zones (i.e. across the stroke). As such, during head frequency assignment, if a first head performs well at ID but poorly at OD, and a second head has the opposite performance, the performance variation between the heads is traded off such that the first head is assigned to high density (frequency) at ID and at low density at OD, and the second head is assigned to low density at ID and high density at OD. In a method of per zone variable bits-per-inch (BPI or linear density) for improving capacity according to the present invention, several manufacturing and customer constraints are taken into consideration. And, performance of each head across the stroke, as well performance variation from one head to another, is utilized in designing the density format and assignments of heads to the density formats.

Referring back to FIG. 1A, an example of density layout according to an embodiment of the present invention is shown. In one aspect, the present invention provides a variable BPI layout as a function of zones on each disk surface 23 based on e.g. two data recording formats (i.e., low density and high density) that utilize: (1) head performance variation from one head to the next in the disk drive and (2) the performance (variation) of a given head across the stroke on a disk surface. Further, the present invention provides a method of generating said layout.

In this example, each zone comprises a group of tracks laid out in zones 60 between one radius and another radius on the disk surface 23, wherein the zone layout for the multiple disk surfaces in each disk drive 100 are the same. A disk drive 100 includes N disks 12 (Disk1 through DiskN), *each disk having a Surface1 and opposing Surface2*, wherein each disk surface has M zones

(Zone1 through ZoneM) across the actuator stroke, and one head per disk surface. For each disk surface, Zone1 is in ID, ZoneM is in OD, wherein the radial boundaries of Zone1 of Surface1 of Disk1 are the same as boundaries of Zone1 of Surface2 of Disk1 *and so on*. Similarly, the radial boundaries of ZoneM on Surface1 of Disk1 are the same as radial boundaries of ZoneM on Surface2 of Disk 1, and so on. However, different zones across the stroke on each disk surface need not necessarily have the same number of tracks or TPI (e.g., Zone1 and ZoneM on the same surface do not necessarily have the same number of tracks). For example, Zone1 *on Surface1 of* Disk1 (i.e., Head1) has same number of tracks and physical zone boundaries as Zone1 on *Surface1 of* DiskN (Head 2xN), etc. And, ZoneM on *Surface1 of* Disk1 (Head1) has the same number of tracks and physical zone boundary as ZoneM on *Surface1 of* DiskN (Head2N). However, the number of tracks in Zone1 and ZoneM can be different. The physical zone boundaries vertically align on the disks in each disk drive, forming virtual cylinders 39 (VC). In this example, there are n virtual cylinders 39, VC1 through VCn. Within a virtual cylinder, different heads may read/write at different frequencies (e.g., variable BPI), hence the concept of vertical zoning.

The level of track density (TPI) can be one of fixed number of preselected levels or can be derived from an algorithm that is based on the location of a portion 35 of the media surface 23. Embedded servo sectors 25 are initially written on a media surface 23 during a factory servo-writing process at a servo track density that can be higher than the data track density, as illustrated in FIG. 1C. Servo bursts within each servo sector 25 are provided in such number and placement to enable accurate positioning of the magnetic head 16 in a full range of positions across the media surface 23, given the particular effective width and characteristics of the read element of a particular head (the read element width typically being narrower then the writer carry out the head positioning method, information in the embedded servo sector 25 is read by the magnetic head 16 and passed to the drive controller 57 which directs the actuator motor 20 to readjust the position the suspension arm 16. In the example shown in FIG. 1C,

the servo track density is about 150% of the maximum possible data track density. In FIG. 1C five servo tracks 37 (e.g., Sa, Sb, Sc, Sd and Se) are shown in relation to three data tracks Tk1, Tk2 and Tk3. Servo track density is determined by determining the max. read and min. write width of a population of magnetic heads 16. After writing the servo wedges 25 at the servo track pitch, the actual data track 30 can be written at any disk radial position between the servo tracks 37. Additional tests, can be performed to determine the optimum data track density of the media surface 23. Each servo track comprises radially similarly situated servo information in servo wedges 25 (e.g., the set of servo information Se at essentially same radial distance from the disk center form a servo track circumferentially, set of servo information Se at essentially same radial distance from the disk center form another servo track circumferentially, etc.)

FIG. 1D shows an example diagram representing the general zone layout of a disk drive with N disks and 2N heads, depicting different heads in a section Zone1 on different disks. FIG. 1E shows another example of capacity zone layout on a surface a disk. FIG. 1F shows example of a series of radial zones, Zone1 through ZoneM, on a disk surface, wherein each zone includes multiple virtual cylinders. FIG. 1G shows an example representative data track layout for each of several virtual cylinders in a zone (e.g., Zone1), on different disk surfaces with corresponding heads. FIG. 1H shows another example servo track 37 and data track 35 layout for a zone 60 on different disk surfaces 23 with corresponding heads, wherein the number of servo tracks and data tracks in different virtual cylinders 39 of a zone (e.g., Zone1) on different disk surfaces 23 are the same. And, FIG. 1I shows another example layout wherein the servo and data track layout varies from zone to zone 60 on a disk surface 23. In all of the above examples of vertical zoning layout according to the present invention, within a virtual cylinder 39, different heads on different surfaces may read/write at different linear frequencies on the data tracks (e.g., variable BPI).

Overview of Format Optimization

In one embodiment, a vertical zoning method according to the present invention includes designing/optimizing (selecting) for two or more recording frequency profiles (i.e., per zone) for a sample number of disk drives (e.g., performed off-line during disk drive development/design phase). Then for a population of disk drives, in each disk drive, each head is assigned to one of the predetermined frequencies for a given zone (e.g., during disk drive manufacturing phase). The assignment step includes assigning a predetermined read/write frequency (BPI) to each head based on a known number of head allocations and the head's performance capability. A head assigned to a higher frequency/density (HD) records more bits on a track, and a head assigned to a lower frequency/density (LD) records less bits on a track.

Referring to the example in FIG. 1A, if the tested performance of Head1 at Zone1 on Surface1 of Disk1 at a given frequency (after full drive read/write and servo calibration/optimization) is better than a desired target performance metric, then that (strong) head, Head1, is considered to have some margin for storing more information than it was originally accounted for. Thus, the designed recording frequency can be increased at Zone1 on Surface1 of Disk1 for Head1 so as to ensure its performance does not fall below the desired target performance metric. If the tested performance of Head2 at Zone1 on Surface2 of Disk1 at that same frequency (after full drive read/write and servo calibration/optimization), is worse than a desired target performance metric, then that (weak) head, Head2, can be compensated for by relaxing the frequency at which Head2 operates, so as to ensure the target performance metric is met. Performing the above tradeoff between the heads for all zones, without loss of overall capacity, provides resulting frequency profiles (i.e. across the stroke) comprising vertically zoned frequency format profiles.

As the above example shows, by compensating for weak Head2, rather than failing the disk drive due to the weak Head2, vertical zoning improves the

disk drive yield. Without application of vertical zoning such a disk drive would not have passed the test limits, and hence would have failed. Furthermore, the format optimizer uses the estimates of performance (i.e., read/write frequency capability) cumulative distribution function at every zone and target performance metric, to design a group of read/write frequency format profiles for weak and strong heads within a given disk drive. The format optimizer also determines the optimal number of, for example, weak versus strong heads.

The format optimizer does not determine which specific head is actually at the lower or higher frequency, but only provides a breakdown of the number of heads at lower frequency and the number of heads at the higher frequency. That breakdown is fixed, performed off-line, and is used during the head assignment process. Then in the assignment process (e.g., during manufacturing test process), out of 2xN heads in a disk drive with N disks, the number of heads that have to be assigned to each predetermined frequency/format is also predetermined (e.g., number heads to assign to low frequency and number of heads to assign to high frequency).

As such, the heads within a set of disk drives are allocated to the predetermined group of read/write frequencies as part of the optimization process to meet capacity and yield requirement for the disk drives. The allocation process allocates a number of heads in a disk drive to one of the predetermined/designed frequencies, however specific assignment of a particular head to a particular frequency is performed as part of the assignment process thereafter. In one example, in a 2 read/write frequency design (high density and low density) for a set of disk drives each with 8 heads, in each disk drive for Zone1 on all disk surfaces, any 3 heads of the 8 heads are allocated to lower frequency and any 5 heads of the 8 heads are allocated to higher frequency in the allocation process based on the performance measurements of all the heads in the set of disk drives. Thereafter, the specific assignment of each particular head to a particular predetermined frequency is performed as part of the

assignment process. For example, in a first disk drive heads 1, 3, 4 are assigned to low frequency, and heads 2, 5, 6, 7, 8 are assigned to high frequency, whereas in a second disk drive heads 2, 3, 8 are assigned to low frequency, and heads 1, 4, 5, 6, 7 are assigned to high frequency, and so on, wherein the specific assignments depend on specific capability of the heads in each disk drive.

The optimal number of heads per frequency (i.e., head allocation) is determined at the same time that the group of read/write frequencies are designed/selected by the format optimizer, by solving a joint constrained optimization problem. For example, in the 8-head disk drive above, for the case of two frequencies (high frequency freq1 and low frequency freq2) at a ratio to a frequency freq, in each vertical zone, allocating 2 heads to freq1 and allocating 6 heads to freq2, provides a first disk drive capacity. Changing said ratio of the frequencies and the number of heads allocated to each frequency, provides a different capacity for that disk drive. As such, the disk drive capacity is a function of the number of heads multiplied by the frequency allocated to each head per zone. For example, if a nominal surface capacity is 1 unit, and if freq1= 4/3 x freq and freq2 = 2/3 x freq, then one head can be at freq1 for every one head at freq2, whereby the average surface capacity is 1 unit. Using the head performance distributions (i.e., head read/write frequency capability distributions at the target performance metric for every zone), the number of heads, the format of the virtual cylinder, and a desired capacity, the format optimizer determines the frequency for each zone in each virtual cylinder and the number of heads in each disk drive allocated to each frequency, in order to achieve that desired capacity. Thereafter, in the assignment process (e.g., as part of a testing of each disk drive), each specific head in a population of disk drives is assigned to one of the predetermined frequencies based on the allocation criteria and specific head performance.

For an example 4-head disk drive, the format optimizer considers 3 heads at freq1 and 1 head at freq2, then 2 heads at freq1 and 2 heads at freq2, and then 1 head at freq1 and 3 heads at freq2. And in each case, using the estimated head performance cumulative distribution functions determines the disk drive yield. The head performance distributions (e.g., BPI distributions) represent percentages of the heads in the disk drives than can operate at different frequency densities. This allows the format optimizer to determine the yield and capacity. In the description of the example embodiment herein, head performance distribution, such as a BPI distribution, represents head frequency capability (probability) cumulative distribution at a target performance metric.

In one version of the optimization process, disk drive yield is maximized while meeting a constraint on capacity. In another version, the capacity is maximized while meeting a constraint on disk drive yield. In the former case, the format optimizer uses a format where the higher frequency is freq1 and the lower frequency is freq2. In the example 4-head disk drive, allocating 2 heads to freq1 and 2 heads to freq2, provides a surface capacity of 1 unit, and capacity of 4 units for the disk drive. If 3 heads are allocated to the higher freq1 and 1 head to the lower frequency freq2, a higher capacity (i.e. 4 and 2/3) is achieved for the disk drive. In that case, the format optimizer lowers the zone recording frequencies to meet that constraint of capacity of 1 unit per surface. As such, for 2 heads at freq1 and 2 heads at freq2, the format optimizer manipulates the difference of those frequencies such that surface capacity always reaches 1 unit, but maximizes yield whereby a maximum number of disk drives qualify and fewest number of disk drives fail to reach required capacity.

As such, according to one embodiment of the present invention, a vertical zoning approach for variable BPI design includes use of an off-line predetermined per zone design of formats based on disk drive data collection and (joint) BPI distribution extraction methods. In one version, a fixed predetermined zone boundary layout is used to design multiple frequency BPI

formats based on representative or actual BPI distributions at one or more desired target metrics (such as off-track error rate), wherein the BPI distributions are extracted from a finite pre-selected set of disk drives.

The collected data is used to extract said BPI distributions for the heads at every (pre-selected) zone, and the per zone design of low/high density formats for the heads is performed offline. The format optimizer solves a constrained joint optimization process off-line to obtain said format designs, using well-known constrained optimization routines. Using joint BPI distributions allows consideration of potential correlation of BPI capability of heads across the stroke as well as individual contribution of each head to the overall drive capacity (or areal density) and yield.

The off-line design of formats allows consideration of other potential constraints that may arise, as additional constraints within the optimizer, and hence solved by the optimizer. For example, as more information is obtained in quantifying the thermal stability constraints of the recording media (which in turn places an upper bound of linear density for the heads) the off-line design provides the ability of not exceeding those limits. If there are data rate limitations in either the write process capability or ASIC component capability, such constraints may be cast within the joint constrained optimizer to ensure said limits are not exceeded.

<u>Overview of Measurement Process</u>

In one implementation of the method of the present invention, a measurement procedure is used to collect data, such that after processing of the collected data, one-dimensional (1-D), two dimensional (2-D) as well as three dimensional (3-D) joint BPI (probability) distributions at a desired read/write target error rate (or any other choice of metric) can be extracted. Data is collected based on head capability measurements taken at different radial positions of the disk. The collected data is used to extract 1-D, 2-D and 3-D

empirical distributions at a target choice of performance metric. The dimensions are dimensions of the distribution, and the distributions represent capability of each head at different radial positions. For example, several disk drives which collectively include 1000 heads are selected for measurement. In a measurement process, record/playback error rate measurements of the 1000 heads from Zone1 to Zone24 of disk surfaces at different frequencies are obtained. Thereafter, in post-measurement processing: (a) the BPI capability of every head at a fixed target metric at e.g. Zone1 is determined in order to obtain a 1-D BPI distribution, (b) the BPI capability of the head at a fixed target metric at e.g. Zone1 and Zone5 is determined in order to obtain a 2-D BPI joint distribution, and (c) the BPI capability of the head at a fixed target metric at e.g. Zone1, Zone5 and Zone20 is determined in order to obtain a 3-D BPI joint distribution.

The BPI distributions are then used as input to the format optimizer to solve three constrained optimization problems to provide head frequency per zone allocations, wherein: (1) one problem maximizes disk drive yield while preserving the same drive capacity, (2) another maximizes the disk drive capacity while preserving the same drive yield and, and (3) another maximizes disk drive yield while ensuring a desired target drive capacity is met at a fixed target track-per-inch (TPI). Additionally, customer related or application specific integrated circuit (ASIC) data rate (limitation) constraints are also utilized. The format optimizer is capable of solving any one of the above-mentioned three problems, wherein one problem can take priority over another depending on the process phase. For example, at earlier development phase of a program where the disk drive components are not matured yet, meeting the drive capacity may be a challenge. In that case, the format optimizer can be used to design variable BPI format profiles by solving the second problem above. As the disk drive components mature such that meeting the drive capacity becomes easier and meeting the drive yield becomes more important, the first may be considered instead. Thereafter, as part of a test process, an algorithm is used to ensure

appropriate head assignment to high/low density (pre-specified) formats per head and per zone or across the head strokes, based on head allocation breakdown of the format optimizer.

Disk drive yield is improved while meeting desired target drive capacity by allowing a frequency format layout (e.g., high and low frequency) with a predetermined number of high and low performing head allocations. Utilizing realistic constraints such as ASIC data rate limitations, the same fixed target TPI is maintained by increasing the average target BPI across the head stroke on a disk surface to achieve the desired disk drive data storage capacity. As such, head performance variation from one head to the next head in the disk drive (and across the stroke across a disk surface) is utilized to allow increasing the areal density of the stored information while preserving the same overall disk drive yield. In one example, a vertical zoning layout method according to the present invention utilizes several design constraints to improve the drive yield using variable low/high BPI design with a fixed predetermined number of head allocations as a function of zones while meeting the target capacity at a fixed target TPI. Head performance variation or correlation across the stroke is also utilized.

Further, the method of present invention takes into consideration the difference in storage capacity of two or mores zones on a disk surface, as it affects overall disk drive capacity. The disk drive capacity is defined as a weighted combination of zone capacities across the stroke on each disk surface. A correlation in the head performance statistics is extracted from one head to another head, and for every given head considered in a set of disks drives across the head stroke on each disk surface.

The joint constrained optimization process determines a per zone target low/high data density format/layout. The optimization process takes into account constraints including customer related requirements such as the requirement of a

minimum logical block count (LBA), monotonic data rate, and maximum data rate requirements at the outer zone areas which can be formulated into (additional) constraints. Head allocation and assignment according to the present invention improves manufacturing yield and provides a disk drive with minimal performance degradation (i.e., sequential or random throughput as well as test process time).

Example Implementation

FIG. 2A shows an example function and flow diagram of an example implementation of the above-described method according to the present invention for generating optimal data density format/layout shown by example in FIG. 1A. The example method in FIG. 2A includes: a data collection/measurement process (block) 62, a data post-measurement-process (block) 64, format optimizer process (block) 66, format generator process and head assignment process (block) 68, example embodiments of which are described below.

Data Collection/Measurement process

In one embodiment, the measurement block 62 implements a measurement procedure including the steps of:

1. Creating several different predetermined linear density format profiles comprising a profile of different frequencies per zone across the actuator stroke, such as e.g. a first profile including freq1_1 for Zone1, freq1_2 for to Zone 2, ..., freq1_M for ZoneM; and second profile including freq2_1 for Zone1, freq2_2 for Zone 2, ..., freq2_M for ZoneM, etc., to be loaded on a representative number of disk drives selected for the measurement process (or if possible on all of the available built drives for that build);

2. Loading a frequency format profile;

3. Performing read/write and servo optimization and calibration;

4. Taking head performance measurements including e.g. (off-track) mean square error (MSE) or quality metric (QM) and/or symbol error rate (SER)

measurements at pre-selected frequencies for preferably all available zones, and saving the data; and

5. Repeating steps 2-4 above for all the remaining frequency format profiles.

The above steps are performed for the selected disk drives in the measurement process.

As such density is selected and data is recorded on a portion of the media surface 23 at the selected density by positioning a magnetic head 16 abutting the portion 35 of the media surface 23, and sending the appropriate signals to the write element (not shown) of the magnetic head 16. Typically, a sample of data such that a significant number of errors are detected (e.g., 10 errors per error rate measurement), is recorded on the media surface 23 to obtain a statistically representative sampling of the error rate for the portion 35 of the media surface 23. Thereafter, the recorded data is read by the read element (not shown) of the magnetic head 16, and the data read is stored by the computer system 54 for evaluation. An error rate of the recorded data is measured or compiled by comparing the actual written data with the read data, element by element. Suitable methods of determining the error rate include actual bit error measurement in which a bit of data read from the media surface 23 is compared with the correct bit, or a correct bit stream in compared with a measured bit stream. An alternative method uses the mean square error metric method in which a waveform read from the media surface 23 is compared with an ideal waveform to provide an error signal that is squared and summed to form the error metric.

In this description, a component distribution is defined to be a (random) variation (i.e., tolerance) of a pre-specified (target) nominal component parameter such as a head write/read width, and the term distribution is defined as the probability distribution function (PDF). During early product development

process, when the head performance distributions are wide and unreliable, data from a matured set of disk drives is used for extracting reference (joint) BPI distributions at a target metric (e.g., off/on track error rate or mean square error). Later in the process, when the amount of head performance variations from one phase to the next in the distribution is expected to be minimal, new sets of measurement data are collected using a selected plurality/population of disk drives at their more matured stages.

Thus, a number of BPI formats including the nominal target format are selected. Then, off or on track error rate or MSE measurements at different pre-selected locations of the disk surfaces, e.g., outer, middle and inner zone are taken (in one example scenario described further below, the choice is limited to said three zones, to reduce the time for performing measurement). However, preferably measurements over multiple zones can be performed and other measurements such as off-track measurements (e.g., 747 measurements) can also be taken. The nominal formats are generated from the data.

Two or more different linear density format profiles can be loaded at a time. In one example, two variable BPI format per zone design (low/high density format profiles) can be created for the purpose of measurement data collection during every build. In this way, more statistical data can be collected from more disk drives, however, there will be only two frequency samples per zone available for data post-measurement processing.

<u>Raw Data post-measurement-process</u>

In the above steps, measurements (e.g. either MSE or SER) for every zone are taken at a finite number of frequency samples. In post-measurement processing (post-processing) block 64, using the available performance metric, measurements are used to calculate each head's frequency performance (e.g., kilo flux per inch (kFCI) or kilo bits per inch (kBPI)) at a given target performance

metric. The performance of every head at every zone is determined as a function of said read/write frequency profiles used for the measurements.

For example, if 6 different frequency profiles are used, then for every head per zone, the measurement process 62 provides measured data as a function of 6 frequency samples at a target metric. In the post-measurement process 64, all the measured data is sorted and performance of every head at every zone at the 6 different frequency samples is extracted to generate frequency capability histograms at a target performance metric (e.g. error rate). Referring to FIG. 2B, the samples can be depicted in a two dimensional graph wherein the x-axis (horizontal) is frequency (e.g., frequencies 1 through 6, kBPI at outer diameter OD), and the y-axis (vertical) is error rate measurement in log scale for Head1 at zone 1 for each of said 6 frequencies (log SER) (in FIG. 2B, each sample data 70 is depicted as a "+"). The error rate measurement can vary e.g. from $1e^{-4}$ (i.e., $1x10^{-4}$) to $1e^{-7}$ (i.e., $1x10^{-7}$) as a function of the 6 frequencies, wherein the error rate increases as the read/write frequency (density) increases.

To determine frequency capability of e.g. Head1 at Zone1 at a target error rate of $1e^{-6}$ (i.e., $1x10^{-6}$), a curve is fit (e.g., using known curve-fitting techniques such as least squares polynomial fit) to the 6 samples, to determine by interpolation the frequency value that gives rise to that target error rate (in FIG. 2B, each curve fit point 72 is depicted by a "o"). If the target error rate is at $1e^{-8}$ (i.e., $1x10^{-8}$) then the frequency value that gives rise to that target error rate is determined by extrapolation (in FIG. 2b, the projected or extrapolated frequency value 74 is shown as a diamond shape). The process for that target error rate is performed for Zone1 for all the heads in the selected disk drives used in the measurement process 62, to create a histogram (FIG. 2D) of the frequency capabilities of all the heads in the disk drives at Zone1 at that fixed target error rate. The process is the same for all zones. As such, using every available head considered in the disk drives under measurement, BPI histograms can be extracted at a given target performance metric per zone.

FIG. 2B shows an example curve of error rate of performance (SER) as a function of BPI (e.g., SER at 6 different BPI/frequency samples) for a head located at the outer diameter (OD) of the disk. Also shown is extracted BPI/frequency capability value of that head at a zone (e.g., OD) for the specified target error rate, using interpolation/extrapolation (i.e., if the specified desired target error rate is outside the performance range, extrapolation or interpolation, such as polynomial fit, is used as necessary). The amount of BPI gain, or margin relative to the nominal BPI setting, is also specified and marked.

The above process is performed for all the heads considered in the measurement procedure, and a BPI/frequency capability of all heads at a given target error rate for every zone is generated. Thus, in this manner, BPI/frequency capability histograms for every zone at a specified target error rate are constructed. If the histogram of BPI capability at a target error rate of every zone is not available, interpolation/extrapolation is preformed to construct histograms for the intermediate zones.

The constructed histograms are used to calculate cumulative performance distribution functions (CDF) of the head frequency (e.g., BPI) capability at a given target error rate performance metric) as input to the format optimizer. Such performance distribution functions are designated as marginal, individual or per zone distributions. In one example the distribution functions include one-dimensional (1-D), two-dimensional (2-D) and three-dimensional (3-D) joint BPI/frequency capability CDF, calculated at the same specified target error rate. Marginal or one-dimensional distribution functions from the joint distribution functions can also be calculated.

A version of estimating performance (e.g., frequency capability) cumulative distribution functions at a given desired target performance metric and zone is described. A number of frequency format profiles can be generated

and tested on a set of disk drives to ensure proper operation. The frequency formats are generated and used to exploit every head's linear density or frequency/BPI sensitivity at every zone. Thus, for example, the linear density sensitivity of every head at a ZoneK (K ranges from 1 to M zones) is determined. To do so, the performance of every head is measured (after full drive read/write and servo calibration/optimization) at each frequency at ZoneK. If Freq1_K, Freq2_K, ..., Freq6_K are the selected frequencies at ZoneK, in the measurement process, every head is positioned on a track (e.g., the same track) at ZoneK and the record/playback performance of each head is measured at every frequency using a performance metric of choice (e.g., off-track symbol error rate (SER)).

For example, FIG. 2B shows the on-track SER performance of a selected head e.g. Head1, at Zone1, and the best least square polynomial fit curve. For a desired on-track target SER of 6e-7, the BPI capability of the (randomly) selected Head1 at that target performance metric can be projected by extrapolating the data (interpolation is performed if the desired on-track target SER is in the performance range, and extrapolation is performed otherwise). FIG. 2B shows the projected BPI capability of that head at an on-track SER of $6e^{-7}$ (i.e., $6x10^{-7}$). The original nominal kBPI (i.e., before the application of vertical zoning) is also shown. The head can be classified as a strong head because there is a reasonably significant amount of margin before the on-track SER performance of this head can be changed from its nominal frequency/kBPI of ~ 188 with an on-track (log of) SER of –9.1 to a projected on-track (log of) SER of –6.22 operating at a frequency/kBPI of ~ 217. Hence, there is a total kBPI gain of ~ 29, allowing increase in the nominal frequency by 15 % while meeting the desired target on-track SER performance metric of 6e-7 (i.e., $6x10^{-7}$). Thus, for example, Head1 of disk drive3 has a frequency capability equal to 217 at an on-track SER of 6e-7 (i.e., $6x10^{-7}$), which is noted for Head1 as one sample for generation and extraction of empirical histograms.

Performing the above steps for all the heads of all the disk drives considered in the measurement process, allows extraction of the empirical histograms of frequency capability at on-track SER of $6e^{-7}$ (i.e., $6x10^{-7}$) for such heads, as shown by example in FIG. 2D. The y-axis shows the number of heads that meet the interpolated/extrapolated (i.e., projected) frequency capability that is shown on the x-axis. The extracted (empirical) histogram can be used to estimate the probability cumulative distribution function. The width of each histogram 76 corresponds to a variance of the head performance histogram, wherein an objective of the present invention is to improve the disk drive yield and capacity, and as a result reduce that variance.

FIG. 2C shows an example joint BPI distribution plot. Such BPI distributions may predict that, for example, 10% of the heads in the disk drives for which measurement was performed, can operate at a frequency density of 1.5 x freq (wherein freq is a reference frequency), and 50% can operate at density of 1.25 x freq, and 90% can operate at density of 1.0 x freq, and 99.9% can operate at density of 0.75 x freq, etc. Using the estimated frequency capability cumulative distribution functions at the target performance metric and every zone, the format optimizer determines the disk drive yield and capacity.

FIG. 2C is an example of a 2-D joint (i.e. outer diameter (OD) and middle diameter (MD)) BPI cumulative distribution function (CDF) at a target performance metric (e.g. on-track symbol error rate of 6e-7). The x-axis shows the BPI capability of all heads (i.e. from all the disk storage devices considered in the measurement phase) at MD AND y-axis is the BPI capability of all heads at OD. The z-axis shows the (calculated) number of heads divided by the total number/population of heads i.e., an estimate of probability that those heads have a joint BPI capability at OD AND MD of less than or equal to any given desired values. BPI capability for e.g. at OD in the above description it is meant that while operating at (or below, i.e., if considering CDF) a given BPI, after full Read/Write and servo calibration and optimization, meeting a desired given

target performance metric of choice at OD. Thus, the aforementioned description of BPI capability can similarly be extended to joint BPI capability.

Format Optimizer process

The format optimizer process block 66 comprises a variable BPI optimization process for solving multiple (e.g., three) constrained optimization problems given the inputs: number of frequency formats (i.e., desired number of different read/write frequencies), number of heads in each disk drive, and said BPI distributions. The first problem maximizes drive yield while preserving the same drive capacity, the second problem maximizes the drive capacity while preserving the same drive yield and the third problem maximizes drive yield while allowing reduced/relaxed TPI such that the desired target drive capacity can be met.

The format optimizer 66 inputs said BPI distributions, a desired disk drive capacity, the total number H of heads per disk drive and the number N of frequency profiles or vertical zones (i.e., frequency per zone across a media surface stroke). Given the frequency capability distribution of heads at a target performance metric, the total number of vertical frequency formats (F), the total number of heads per drive (H) and the nominal drive capacity, the format optimizer 66 simultaneously searches through all possible continuous range of frequency capabilities to maximize the drive yield such that the desired nominal drive capacity is met. The format optimizer 66 can also solve the same problem, but the drive capacity and drive yield interchanged.

As such, in one example, the format optimizer 66 optimizes high/low density as a function of zone, to improve the drive yield and meet a fixed target capacity. The format optimizer also optimizes capacity while achieving a fixed nominal drive yield, wherein the nominal yield is before the application of vertical zoning according to the present invention. In said example of F=2 formats (high density and low density) and H=8 heads 16 in a disk drive 100, the possibilities

include 1 head at high density and 7 heads at low density, 2 heads at high density and 6 heads at low density, 3 heads at high density and 5 heads at low density, 4 heads at high density and 4 heads at low density, 1 head at low density and 7 heads at high density, 2 heads at low density and 6 heads at high density, 3 heads at low density and 5 heads at high density, etc. Hence, the format optimizer considers all the combinatorial possibilities, and in each case solves a constrained optimization problem and finally chooses the best optimal solution amongst all the possibilities. Alternatively, the format optimizer can be designed to reach the best optimal solution more directly by solving a (non-linear) mixed integer programming.

Therefore, once the 1-D, 2-D and 3-D joint BPI/frequency capability CDFs (discussed above) at a given target performance metric are calculated and passed as input to the format optimizer, the format optimizer solves the above two problems, namely: (1) maximizing or improving the drive yield (i.e., due to the pre-selected performance metric, e.g., off-track SER) while meeting a desired nominal drive capacity and (2) maximizing the disk drive capacity while meeting a desired nominal drive yield.

The format optimizer then mathematically casts the two problems stated above as constrained optimization problems and solves them using well-known optimization techniques such as e.g. line search algorithm. The constrained optimization problems can also be cast as (non-linear) mixed-integer programming and solved using existing methods in optimization theory. Example constraints to be considered, and cast mathematically within the format optimizer, include not exceeding a certain frequency at OD due to ASIC data rate limitations or at ID due to head/media limitations. Furthermore, closed form equations are derived and used in the format optimizer to estimate the actual drive yield and drive capacity. A Format Generator process, described below, is utilized to calculate the actual drive capacity after including all possible

overheads, such as adding and including redundant bits due to error correction coding or gray coding.

The format optimizer uses the information from the format generator, such as the calculated format efficiency per zone (i.e., defined in percentages as the amount of user data e.g. in blocks that can fit in all tracks in a zone), or number of tracks per zone, to achieve a very close estimate of disk drive capacity calculation as determined by the format generator. Then, the format optimizer 66 calculates optimal linear density format profiles as well as optimal number of heads allocated to each vertically zoned format profile.

As an example, histograms are extracted and corresponding distributions are estimated for different zones at desired target error rates (e.g., for Zone1 at a target error rate of $1e^{-6}$, Zone2 at a target error rate of $1e^{-6}$, Zone3 at a target error rate of $1e^{-6}$, etc.) as described above. A format design is provided for 4-head disk drives (H=4), and 2 vertical frequency format/profiles (F=2), wherein the disk drive yield is optimized while meeting a minimum capacity requirement.

Without the method of the present invention (vertical zoning), conventionally when the same frequency is used per head per zone, if one of the 4 heads is a weak performing head having an error rate measurement e.g. $1e^{-5}$ at Zone1 (higher than the target error rate), that disk drive is failed. With the application of the present invention in that case, the format optimizer allocates the 3 other heads to higher frequencies, and allocates the weak head to a lower frequency at Zone1. The recording/playback performance of the weak head is compensated for, such that the minimum capacity requirement is met. As such, the format optimizer 66 utilizes the performance distributions to determine two or more optimal frequencies per zone, and the optimal number of head allocations to those frequencies per zone such that constraints such as required disk drive yield and/or capacity are met.

For example, the format optimizer uses the estimated 1D, 2D and 3D joint frequency/BPI capability distributions at a desired target performance metric to jointly optimize for vertically zoned frequency format profiles and the corresponding number of head allocations three zones at a time. An advantage of considering three zones, as compared to only one (thus considering a joint optimization versus individualized optimization), is that joint optimization allows optimization of format profiles (e.g., frequency profiles) across the stroke on each disk surface. Therefore, in this way we exploit the potential correlation in performance from one zone to another as well as their individual and weighted contribution to the overall surface capacity. A joint optimization approach is preferable for a low/high density format layout across the stroke for either improving the drive yield while keeping the same drive capacity, or improving the drive capacity while preserving the same drive yield.

The generated results from the format optimizer include: the target high/low BPI formats per zone, (optimal) number of head allocations per format/layout and an estimate of the drive yield and capacity. The accuracy of the estimates can be sensitive to the underlying extracted (joint) BPI distributions at a given (on or off-track) error rate. Further, the target high/low BPI formats can be sensitive to the variance of the (extracted) BPI distributions. And, the variance of the BPI distribution can be sensitive to the absolute value of the (on or off-track) target error rate and the choice of metric (e.g. error rate vs. MSE). In addition, because the design of target high/low BPI formats are performed three zones at a time and the yield improvement (while preserving the same overall drive capacity) is based on the profile of the target nominal formats, the format optimizer allows for a smoothing operation in order to smooth the target variable BPI format designs if so desired. The format generator determines the number of tracks per zone, the number of blocks per track, the radius at each zone, as well as block and track format efficiency. This information is saved in e.g. output files for use with the format optimizer. The format optimizer then saves the

design of target high/low BPI formats per zone that it generates, in two separate files that can be read and loaded as input files into the format generator.

Once the target format profiles are calculated, if they are non-smooth across the stroke, optionally a smoothing process is applied. The format profiles are then loaded into the format generator 68 described below, to create vertically zoned formats and configuration pages. The formats and configuration pages are used by the disk drive firmware to create binary files to be loaded onto the reserve image of the drives as part of the file system. FIG. 2A shows the communication between the format optimizer 66 and the format generator 68. In this fashion the design and implementation of format profiles, as well as the number of optimal head allocations are performed off-line and are predetermined for every drive configuration.

An example of format optimization for designing vertically zoned low and high frequency profiles for disk drives with 4 heads (i.e., Head1 through Head4 corresponding to Surface1 through Surface 4 of Disk1 and Disk2) is described. In this example, every disk surface is (uniformly) partitioned into three zones across the stroke, with a fixed number of tracks (TPI) per zone, vertically aligned from one surface to another. The nominal surface capacity (before the application of vertical zoning) can be approximated by the sum over all zones of the products of nominal tracks per zone by the nominal frequency per zone and by the format efficiency per zone. Format efficiency per zone is a percentage of all the user data that is effectively stored per zone. Then the nominal disk drive capacity is equal to the nominal surface capacity multiplied by the total number of surfaces (or heads). The nominal number of tracks per zone and the format efficiency per zone can be generated and obtained from the format generator (described further below).

Performing vertical zoning to e.g. improve the drive yield without losing any nominal (disk drive) capacity, finds the best frequency per zone and per

head such that the disk drive meets performance and capacity requirements. As such, if a disk drive with 4 heads fails test process performance limits due to e.g. performance of Head1 at Zone1, but Head1 at Zone2 (or another head such as Head3 at Zone1) performance is significantly better (i.e., passing the test limits with reasonable margins), then a higher frequency than nominal at Zone2 or Zone1 is designed for heads that are stronger (i.e. high density heads), and instead the frequency at Zone1 for heads that are weaker (i.e. low density heads) is lowered. This tradeoff is performed such that the overall disk drive capacity is preserved, to obtain a vertically zoned design of variable frequencies per zone. In addition, the number of heads (e.g. per zone) allocated to high or low density is determined. Thus, the disk drive capacity can be approximated by the sum (over all zones) of the products of number of low density heads per zone by low frequency per zone by format efficiency per zone plus the (sum over all zones of the) products of number of high density heads per zone by the number of high frequency heads per zone by the format efficiency per zone.

In one version, the format optimizer solves for the above problem as follows. The format optimizer is provided with the (joint) frequency capability cumulative distribution functions (extracted and estimated from all heads considered in the measurement process above) at a desired target performance metric (i.e., the same targets used in the test process). Then, for every combinatorial possibility of head allocation (e.g. to high or low frequency) the format optimizer searches through a continuous range of possible frequencies, by considering every zone independently (i.e. using the marginal distributions) and by combination of zones (i.e. using the joint distributions), to maximize the disk drive yield calculated using a closed form equation, such that the disk drive capacity after the application of vertical zoning is essentially the same as the nominal disk drive capacity. Further, the (optimal) low and high frequency profiles for every combination of head allocations is compared and the one that results in the highest value of (calculated) disk drive yield is chosen and passed

to the format generator for the generation of vertically zoned configuration pages to be used by the disk drive firmware.

The surfaces of disks can be partitioned to more than the example three zones above. The above steps of determining the optimal variable frequencies per zones are useful to consider more than three zones. To reduce computational complexity and time, if the selected/designed number of zones per surface in a disk drive is more than three, the format optimizer can be used to generate low and high frequency profiles three zones at a time and, suitable smoothing operations are used to smooth the profile after post-processing. Another approach includes embedding the smoothing operator in the design and extend the joint optimization to all zones so as to consider the effect and impact of smoothing to drive yield (calculation) as part of design rather than later stages.

In the above 4-head disk drive example, disk drive yield is maximized while preserving the same nominal disk drive capacity. To determine the number of head allocations, the format optimizer begins with one low density head and three high density heads (e.g., per zone). The format optimizer searches through a continuous range of possible frequency capabilities per zone, as well as two and three zones at a time, by considering and using the 1D (i.e., marginal), (joint) 2D and 3D distributions that result in the best calculated value of the drive yield such that a minimum nominal drive capacity can be obtained. Next, the format optimizer uses two low and high-density heads and repeats solving the same constrained optimization problem. This process is continued until all the combinatorial possibilities are considered. Finally, the format optimizer chooses the solution that results in the best calculated value of the drive yield and provides the target low and high (optimal) density format profiles and the associated number of high and low head allocations to the format generator. The format generator then generates vertically zoned format files and configuration pages to be used by the firmware.

Format Generator

In one embodiment, the format generator process block 68 is used for embedded servoing (i.e., servo position is generated by reading back written information from the disks, such as servo wedges in which position information is embedded on the disks, and that information is used to position the head on the disk surfaces).

For example, the format generator 68 described herein generally performs three functions including utilizing target formats/frequencies (or linear densities/BPI) for each zone as an input, and performing an exact calculation of the capacity of each zone and the disk drive itself. Further, the format generator 68 calculates the format efficiency (percent of the disk area that is occupied by customer data) for each zone. The third, and primary purpose of the format generator 68 is to generate configuration page data. The configuration pages contain per-drive, per-zone, and per-head-per-zone parameters that are programmed into disk drive electronics components. Such components include the disk controller, the read/channel, and the preamplifier. The parameters are ordered such that the disk drive firmware selects the correct set of parameters to be programmed into each of the components for the particular head and zone that is being written to or read from at the time.

The format generator 68 calculates the exact frequency and the exact capacity of each zone taking into consideration limitations in the programmability of the components and limitations of the capabilities of disk drive components. Some examples of component limitations include: read channel synthesizer frequencies are limited to a discrete collection of frequencies; preamplifier (not shown) has a minimum and a maximum delay in turning on its write current; the down track separation between the head read and write elements (not shown) varies from component to component; the reference crystal (not shown) has finite accuracy and stability over temperature; the spindle motor driver can keep the spindle motor speed within a finite precision of the nominal rotational speed;

the controller has specific latencies in generating commands to the read channel and preamplifier, often with a finite uncertainty as to the exact timing of these commands, etc.

The format generator 68 can be fully automated, or can be directed by a human specialist. In the absence of input from the format optimizer 66, the target per-zone BPI/frequency profiles, in particular, must be generated from human input. In general, the human specialist modifies the target frequency profiles until the desired capacity is reached.

In one embodiment, the format generator 68 includes a format efficiency process that uses the format optimizer's target low/high variable BPI format designs as well as the optimal predetermined number of low/high performing head allocations, to modify and generate the appropriate configuration pages (i.e., as part of the file system). For each zone, the format generator 68 selects the nearest frequency to the target frequency for that zone, given the component limitations and programmability mentioned above. The nearest frequency comprises the target formats.

The optimal predetermined number of low/high performing head allocations comprises the number of heads allocated to each of the multiple frequencies in each zone. The format optimizer 66 determines the head allocation, which is input to the format generator 68. The capacity of a zone depends both on the target frequencies and the number of heads allocated to each frequency.

The format optimizer uses the nominal average BPI or frequency (nominal BPI format target designs) (e.g., one read/write frequency) in each zone as input from the format generator 68 to estimate the disk drive yield before applying variable BPI designs. For a design with multiple frequencies per zone, this is the weighted average (by the number of allocated heads) of the multiple frequencies.

The nominal format is created by e.g. a human operator working with the format generator 68 in the interactive manner described above.

The format generator 68 performs calculations of number of tracks per zone, number of blocks per track, radius at each zone as well as block and track format efficiency to calculate the drive zone capacity. The format optimizer 66 estimates the capacity using the tracks per zone, radii, and format efficiency. Thus, the format optimizer and the format generator interact as shown in FIG. 6. For example, for 4-head disk drives, and 2 density format frequency profiles (i.e., high and low frequency profiles) with 3 zones across the disk surface, after the measurement and optimization processes, the format generator 68 is provided with 2 optimal frequency profiles and optimal allocation of the heads. The format generator 68 then calculates capacity, and if the drive capacity meets the minimum required capacity, then the format generator generates configuration files/pages for the disk drive firmware. The configuration pages are used by the drive firmware to command the head to write at an assigned frequency to a zone. If the calculated drive capacity does not meet the minimum required capacity, format optimization is performed again with new format efficiency values, and the process is repeated.

<u>Head format assignment and selection criterion</u>

The process for allocation of numbers of heads to each of the predetermined multiple frequencies in a zone, and the process of assignment of a particular head in a particular disk drive to a frequency, are distinct. First the allocation process is performed by the format optimizer 66, and applies to disk drives of a particular design (product). Then, the assignment process is performed during manufacturing as part of a test process undergone by each disk drive to be produced. This section describes the assignment process task.

Once the configuration pages are generated and converted to binary files as part of the file system, they can be loaded and saved in a reserved image of

the disk drive for use after power cycling. Then, for every disk drive, the following example assignment process is performed per head and per zone, to determine assignment of a certain predetermined number of heads to high BPI formats and the remaining heads to low BPI formats in a 2 frequency design, to satisfy allocation of heads to said formats by the format optimizer 66.

The assignment process for the example 2 frequency format where high and low frequencies are used, includes the steps of:

1. Load default parameters from the configuration pages, and calibrate selected parameters on a per head, per zone basis (e.g., load high BPI format profile for all zones across the stroke);

2. Take measurements from all heads at all the disk surfaces at pre-selected zones with respect to a metric, e.g., mean square error, on/off track error rate, etc.;

3. For each head in every measured zone, sort/rank the heads by the performance metric from best to worst; select a pre-specified (by the allocation process in the format optimizer 66) number of heads with the best performance, and assign those heads to the higher frequency for a particular zone;

4. Optionally interpolate between the measurements obtained from the pre-selected number of zones to find the results for the other zones, and do the same for the interpolated zones. The interpolation operation in a version of the present invention reduces test process time. Head performances are measured, sorted and assigned to a frequency for a subset of the total number of zones. For the remaining zones, the heads are assigned by interpolating on the head assignments made from measurements;

5. For every zone, save the worst pre-specified number of bad (i.e. low performing) heads with respect to the selected metric; and

6. For every zone, load and calibrate all the bad heads with the lower BPI format.

The above process can be used to improve yield, improve capacity, and trade off between yield and capacity. In a test, the heads can be passed or failed with respect to a metric (e.g., off track error rate) to determine if the test target limits are met.

The servo firmware is extended to load more than one format profile. A head can be assigned a different read/write frequency per zone across a disk surface, and radially similarly situated zones on different disk surfaces can have different read/write frequencies assigned to the corresponding heads whereby one head is assigned a different frequency/format profile than another head.

The example assignment process described herein applies to a format design with two recording frequencies per zone. However, the process can be easily extended to more than two frequencies per zone, wherein the process can be iterated upon to assign heads to more than two frequencies per zone, as described by an example below. For a design with H heads and F frequencies per zone, above steps 1 and 2 are completed for the high frequency. The first selection of heads in step 3 assigns the highest h_1 heads, where h_1 is the pre-specified number of heads allocated to the highest frequency for that zone. The remaining (H - h_1) heads are then loaded and calibrated with the second highest frequency (step 1 again), measurements taken (step 2 again), and the heads ordered relative to the metric and the best h_2 heads are assigned to the second highest frequency (step 3 again). Here h_2 is the pre-specified number of heads allocated to the second highest frequency in the zone. Steps 1-3 of the process are then iterated for the (H - h_1 - h_2) heads, followed by the (H - h_1 - h_2 - h_3) heads, and so on, until h_F heads remain to be assigned to the lowest frequency. The set of {h_1, ..., h_F} heads comprise the head allocation made by the format optimizer 66.

Table 1 below illustrates the result of an example of the vertical zoning head assignment process on a disk drive with 6 heads and 5 zones across the

stroke on each disk surface. Each head is assigned to either high or low density format based on record/playback performance of that head, wherein as discussed above, the number of heads assigned to low density and the number of heads assigned to high density is according to the head allocation results determined by the format optimizer.

HEAD #	FORMAT\ZONE ZONE 1	ZONE 2	ZONE 3	ZONE 4	ZONE 5
0	Low	High	Low	High	Low
1	High	Low	High	High	Low
2	High	Low	High	Low	High
3	High	High	Low	High	High
4	Low	High	High	High	High
5	High	High	High	Low	High

Table 1 – An example for the format assignment of a disk drive after test process, using vertical zoning with variable BPI across the zones.

FIGs. 3-6 show example steps of an embodiment of the above processes. Referring to FIG. 3, an example vertical zoning data collection process includes the steps of:

(1) Select a number of disk drives for data measurement/collection process (step 300);

(2) Create nominal linear density profile KFCI (i.e., nominal_KFCI): $\overline{\mathrm{kFCI}}(R)$, wherein R is disk radius (step 302);

(3) Create more linear density profiles by multiplying the nominal_KFCI

$$(1 \pm x_i) * \overline{\mathrm{kFCI}}(R)$$

by scaling factor x_i, (step 304):

(4) Create binary file system for every generated profile (step 306), i.e. for

$$i \in \{1, \Lambda, N\}$$

wherein N is the total number of frequency format profiles. For example, for N=2, having X_1, and X_2, if X_1=0.05 and X_2=0.1, then including the nominal frequency format itself, there are 5 different frequency profiles in step 304, as follows: (a) nominal_KFCI, (b) 1.05 * nominal_KFCI, (c) 0.95 * nominal_KFCI, (d) 1.1 * nominal_KFCI, and (e) 0.90 * nominal_KFCI (wherein "*" is the multiplication operator);

(5) Select the first head by setting *i* to 1(step 308);

(6) Load file system *i* onto the reserved image of the disk drives (step 310);

(7) Take head performance measurements (e.g., (on/off) track MSE and SER measurements) (step 312);

(8) Unload and save the results in the Data Base (step 314);

(9) Increment *i* by one (*i*=*i*+1) (step 316);

(10) is *i* = N? (step 318)

If not, go to step 310, else done.

The above process collects performance data for all the heads at all zones.

Referring to FIG. 4, an example vertical zoning post-measurement processing and per zone BPI distribution extraction process includes the steps of:

(1) Organize the performance data (e.g., MSE and SER) obtained above, for every head $i \in \{1,\Lambda ,M_1\}$ and every zone $j \in \{1,\Lambda ,M_2\}$ as a function of linear density samples, wherein in this example M_1 is the total number of heads in the disk drives selected for the measurement process (e.g., 40 disk drives each including 4 heads, results in total of M1=160 heads), and M_2 is the total number of zones, to generate head performance histograms (step 400);

(2) Choose a target performance metric (e.g., MSE or SER) (step 402);

(3) Set $j = 1$ and $i = 1$ (step 404);

(4) Interpolate/Extrapolate BPI at the specified target performance metric (e.g., MSE or SER) for head i at zone j (step 406);

(5) Select the next head by incrementing i by one (i = i +1) (step 408);

(6) Is $i = M_1$? (i.e., have all heads been processed?) (Step 410);

(7) If not, got to step 406 to process the next head, else generate frequency capability histogram at that given zone j for all the heads (step 412);

(8) Is $j = M_2$? (i.e., have all zones been processed?) (Step 414);

(9) If yes, stop;

(10) Otherwise, move to the next zone and start with the first head again, whereby $j = j + 1$, $i = 1$ (step 416), and go to step 406 to repeat.

The above process generates (1D) frequency capability histograms at a target performance metric and for every zone by considering all the heads from the sample disk drives selected in the measurement process. Using the (1D) frequency capability histograms at a given target performance metric, techniques in probability theory known to those skilled in the art can be adopted to estimate the (1D) frequency capability distributions. Further, the above procedure above is extended (i.e., by using 2D and 3D interpolation/extrapolation routines), to extract and estimate the 2D and 3D joint frequency capability histograms and their associated distributions.

Referring to FIG. 5, an example head assignment process for a two frequency format (N=2, high/low density) design, includes the steps of:

(1) Assign all n heads in a disk drive to the first selected format (e.g., high-density format) (step 500);

(2) Calibrate all n heads at high-density format for selected zones (step 502);

(3) Measure performance metric at selected zones for all n heads (step 504);

(4) For each selected zone, rank heads by metric (step 506);

(5) For each selected zone, assign highest K heads to high-density format, other n-K heads to low density (step 508);

(6) Optionally interpolate head assignment for remaining zones (step 510); and

(7) Complete calibration of all heads and all zones at assigned formats (step 512).

The above process completes assignment of each head in each disk drive to a predetermined frequency.

As shown in FIG. 2A, information is passed between the format generator 68 and the format optimizer 66, wherein initially, the format generator 68 passes information including e.g. track layout to the format optimizer 66 to have a more accurate way of calculating capacity (nominal format). Such information and constraints are provided to the format optimizer 66 to solve said joint optimization problems. The format optimizer 66 performs a coarse calculation of capacity, whereas the format generator 68 performs an exact calculation of capacity. The format generator performs functions of providing format information (e.g., number of tracks per zone, and zone layout) to the format optimizer 66, and calculating exact format capacity. Such information is passed once from the format generator 68 to the format optimizer 66 for a design (e.g., 4 head design). The format generator 68 initially provides nominal information to the format optimizer 66, wherein the format optimizer 66 performs its calculation of target densities (zone frequencies and number of heads allocated to each frequency) and provides that information to the Format Generator. The format generator 68 then determines if required capacity has been reached. Adjusting target densities to meet yield and/or capacity requirements includes adjusting selected zone density or zone frequencies.

Referring to FIG. 6, an example Format Generator/Format Optimizer Iteration process for minimum capacity requirement C, and a user specified allowed overcapacity Delta, includes the steps of:

(1) Determine disk geometry, track density (TPI) and servo wedge details, and output inner diameter (ID) and outer diameter (OD) radii; TPI; number of servo wedges, and wedge length (step 600);

(2) format generator 68 generates initial format at capacity using the values ID, OD radii, TPI, number of wedges and wedge length, and outputs radius of each zone per disk surface, number of tracks per zone, number of blocks per track, and format efficiency by zone (step 602);

(3) format optimizer 66 generates optimal target densities at all zones as described above, and outputs frequency density targets (e.g., low/high BPI) by zone, and number of frequency density (e.g., low/high BPI) head allocations by zone (step 604);

(4) format generator 68 generates new formats with Capacity (i.e., number of logical blocks per disk drive) (step 606);

(5) If Capacity > C and Capacity < (C + Delta) (step 608), then stop;

(6) Otherwise, adjust target densities (Step 610), and go to step 606.

In one example, surface capacity is described by the equation: TPI x BPI x (1 + ECC) / FE, wherein TPI is track density, BPI is the linear density, ECC is the fractional level of error correcting code used which is typically about 0.1, and FE is the format efficiency which is typically about 0.57.

The above process completes the format generation process.

As an example scenario of the results generated by an embodiment of vertical zoning according to the present invention, a set of 32 matured disk drives are selected wherein each disk drive includes 12 heads. The 1D, 2D and 3D joint BPI empirical distributions are extracted at a given specified target on track error rate from three pre-specified radial zones, i.e., outer, middle and inner

zones. Next, the extracted distributions are fed into the format optimizer, and low or high frequency per zone format designs obtained at the three specified zones. This is performed once by individual optimization, all based on 1D distributions at each of the three zones, and once by joint optimization based on the measurements obtained from the three zones and their extracted 1D, 2D and 3D distributions. The head format allocation search process (VZ test) is performed in a simulation, wherein for each zone the one-format designs (i.e. before the application of vertical zoning) are a special case of two-format variable BPI designs by forcing low and high formats to be the same and equal to the nominal BPI format at that zone. Furthermore, the pass/fail of disk drives is decided based on the criterion that each head at every zone pass a given on-track target error rate as well as off-track squeeze and un-squeeze offset margins. Then, the drive yield is calculated (i.e., in simulation by interpolation/extrapolation of the measurement data) before and after the application of vertical zoning. The following Table 2 summarizes the results:

	Using Joint Optimization	Using Individual Optimization
Drive Yield (Yd)	93.75	90.625
Drives failed after VZ	4 & 29	4, 6 & 29
Drives recovered	2, 3, 13, 19, 21 & 25	2, 3, 13, 19, 21 & 25
Passed drives failed after VZ	None	6
Drives failed before VZ	2, 3, 4, 12, 19, 21 & 29 i.e., drive yield before VZ Yd=75%	2, 3, 4, 12, 19, 21 & 29 i.e., Drive yield after VZ Yd = 75%

Table 2

In addition to disk drives, the present invention is useful with other storage devices such as e.g. tape drives, optical drives, etc. Though a manufacturing test case for a two format design is described, the search algorithm can easily be generalized to higher number of formats. The design of two formats based on 1D, 2D and 3D joint storage density BPI distributions can easily be generalized to higher order or dimensions by considering more zones than three. The design of formats can be generalized from two to higher number of formats. The measurement procedure can be generalized to consider more zones as well as off track measurements such as 747 curves or quality metric versus error rate

measurements to perform correlation study for the choice of best metric with less potential test time.

Further, the above methods for a per zone variable BPI design can be easily extended to a variable BPI/TPI design as described below. The measurement process is extended to further include 747 measurements of all the heads from a pre-selected number of disk drives. To speedup the measurement of raw data, instead of 747 measurements, off-track and adjacency margin (squeeze measurements) of all heads can be performed. Once the 747 raw data of all heads at pre-selected number of zone locations is determined, for every zone (joint) BPI/TPI distributions can be extracted at given desired target(s) by post-measurement processing of data. The choice of target is an integral part of the amount of performance gain, such as disk drive yield, due to per zone variable BPI/TPI designs. Some example choices of target(s) are off-track error rate, the variance of position error signal or even a combination of both. After the joint BPI/TPI distributions are extracted and available for all zones, a per zone variable BPI/TPI design can be obtained by solving two-constrained (joint) optimization problems: one that maximizes the drive yield while keeping the same drive areal density and another that maximizes the drive areal density while keeping the same drive yield. Once the per zone variable BPI/TPI designs are obtained, a head BPI/TPI allocation and selection criterion, similar to that described herein, can be used such that a pre-selected number of heads are allocated to high and low density BPI and TPI formats, for example, for the case of two variable BPI/TPI per zone design performed as part of test process.

The present invention improves drive yield and drive capacity (or consequently areal density at a fixed target BPI), and allows reducing the target TPI by increasing the average BPI across the stroke per head (depending on the number of formats considered) to meet a desire target drive capacity. In particular, due to a maximum deliverable data rate of the ASIC components (e.g., channel, controller and preamp) the BPI at the outer diameter may be limited by

the minimum deliverable data rate of the mentioned ASIC components. For example, if the controller is capable of maximum deliverable data rate of 650 MHz, the preamp capable of 700 MHz and the channel capable of 750 MHz, then the BPI at the outer diameter is limited by the controller at a maximum deliverable data rate of 650 MHz. Thus, a conventional one format BPI profile across the stroke does not achieve the desired drive capacity and the desired manufacturing drive yield. The target BPI is increased and the BPI profile across the stroke is relaxed, wherein according to the present invention, the per zone variable BPI design can be used to design (variable BPI) target formats that meet the desired drive capacity at a fixed target TPI while improving the overall drive yield.

Referring back to FIG. 2A and FIGs. 3-6, in one embodiment of the present invention, the steps of the example method of the present invention: (1) the data collection/measurement process block 62, can be implemented on general purpose computing equipment 61, known in the art, and drive electronics including special purpose electronic circuit (e.g., logic circuit) 49 and on board microprocessor 57 (FIG. 1B), configured according to the present invention, wherein the special purpose electronic circuit 49 is configured to perform the measurements, the on board microprocessor 57 directs the special purpose circuit 49, and transfers the data to the general purpose computer 61, (2) the head assignment process can be implemented on the on board microprocessor 57 within the disk drive 100 configured according to the present invention, wherein a data collection subtask is related to the head assignment task such that the data collection sub-task is performed by the special purpose electronic circuit 49 within the disk drive 100, (3) the steps in each of the data processing block 64, the format optimizer block 66 and the format generator block 68 can be implemented on general purpose computing equipment 61 (e.g., high end PC, PC server or workstation, etc., including programmable simulation software) configured according to the present invention.

The present invention has been described in considerable detail with reference to certain preferred versions thereof; however, other versions are possible. Therefore, the spirit and scope of the appended claims should not be limited to the description of the preferred versions contained herein.